

APPLICATION  
FOR  
UNITED STATES LETTERS PATENT

TITLE: TRANSFER OF MESSAGES IN A MULTIPLEXED  
SYSTEM

APPLICANT: ANTONIO CANTONI AND ROBERT M. NEWMAN

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL224699579US

I hereby certify under 37 CFR §1.10 that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit

Signature

Typed or Printed Name of Person Signing Certificate

July 31, 2001

Samantha Bell

Samantha Bell

10699-003003

## TRANSFER OF MESSAGES IN A MULTIPLEXED SYSTEM

This invention relates to a transfer of messages in a multiplexed system.

More particularly, the invention relates to a method for the transfer of messages in a time multiplexed slotted environment such as a communications network. The network can be of the type disclosed in International Publication No. WO 86/03639 and that disclosure is incorporated herein by cross-reference.

Generally speaking, the invention provides an efficient method for the connectionless or connection oriented transfer of messages of arbitrary but finite length in a time multiplexed slotted environment with constrained destination resources.

The method provides for the efficient support of any type of addressing (short or extended, hierarchical or non-hierarchical) in the one environment, even in a system with short slots.

In the one environment, the method can guarantee delivery of messages or provide a more efficient transfer at the expense of occasional message loss. In either case, there are minimal communication overheads and the utilisation of destination resources is maximised. Thus the method provides a wide range of options and considerable scope for achieving a range of performance-cost objectives.

### BACKGROUND ART

In the data communications environment, information is generally exchanged in units called packets. These consist of an overhead necessary for the control and addressing of the unit through the data switch and of the actual information. Typically the size of the information unit is not fixed but depends upon the message and the amount of information to be transferred.

Early packet switches handled the variable length packets as a whole unit, allocating all its communication resource to the transfer of the packet until its completion. There are a number of new packet switch designs emerging that switch only small fixed length slots. These switches are commonly referred to as Fast Packet Switches. Such switches are an improvement since they are generally simpler, may operate at higher speeds and allow for the support of real-time traffic.

If the Fast Packet Switch is to carry packet communications of variable length then it is necessary that the original packet message be segmented for transmission over the switch and reassembled at the destination. The segmentation function is relatively simple only requiring that the message be divided into units of size equal to or smaller than the slot size. The transmission of the segments over the packet switch however requires much more since it is necessary that the destination can receive and order all of the segments of the message. Thus it is required that there be a logical association between all of the slots of a single message. The reassembly function then reconstructs the original variable length message from all of the received segments. Special care may need to be taken in the reassembly function to allow for the possibility that more than one message may need to be reassembled concurrently.

Some schemes have been developed to perform the segmentation and reassembly function. However, these are limited either in the efficiency of the transport in the switch or in the performance of the reassembly func-

09916725-073101

The problem with the slot overhead is compounded by addressing requirements. The common address field sizes used in data communications are 16 to 48 bits. With 48 bit addressing there is an overhead of 12 bytes per slot (source and destination address) in addition to the reassembly overheads. This approach to segmentation is clearly inefficient with small (less than 32 bytes) slot sizes.

40 BRIEF DESCRIPTION OF THE INVENTION

The invention also provides an apparatus for transmitting variable length messages on a network from a source to a destination in fixed length slots said apparatus including:

a segmentation machine for segmenting the message into fixed length slots which include a header field and a message segment, said machine including coding means for providing a source identifier field in the header of each slot, said source identifier field including a source identifier code which is uniquely associated with the message to be transmitted, and a reassembly machine located, in use, at the destination, said reassembly machine including control means for controlling reassembly of slots in accordance with the source identifier codes of the slots.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagrammatic representation showing a variable length message and fixed length segments;

FIG. 2 is a block diagram showing the segmentation and reassembly machines coupled to a fast packet switch;

FIG. 3 diagrammatically illustrates the segmentation and reassembly machines coupled to a QPSX network;

FIG. 4 is a block diagram of an access unit (AU) of the network shown in FIG. 3;

FIG. 5 is a more detailed block diagram of the access unit;

FIG. 6 is a segmentation state machine diagram;

FIG. 7 is a diagram showing request and confirm primitives for segment transfer to the access unit;

FIG. 8 is a reassembly state machine diagram;

FIG. 9 is a diagram showing an indication primitive for segment transfer from an access unit;

FIG. 10 is another reassembly state machine diagram;

FIG. 11 is a block diagram of a reassembly machine; and FIGS. 12A and 12B are flowcharts illustrating logical operations in the reassembly machine.

## DETAILED DESCRIPTION OF THE EMBODIMENTS OF THE INVENTION

FIG. 1 diagrammatically shows a message which is of variable length. The message includes address fields 22 and 24 for the destination address (DA) and source address (SA). The message includes a length field (L) 26 and an information field (IU) 28 followed by an error checking field (CRC) 30. The error checking field 30 can be of any known type.

In accordance with the invention, the variable length message 20 is segmented into a number of slots 32 of equal length. The method places the address field 22 and 24 in the first of the slots 32 and subsequently logically associates the following slots of the message with the first slot using a unique identifier, as is diagrammatically illustrated in FIG. 1.

Each slot 32 has an Access Control Field (ACF) 34, TYPE field 36 and Source Identifier (SI) field 38. In accordance with the invention, the SI field 38 is used to provide the logical linking between slots 32 of the same message. The slots 32 also include an information field 40 which is used to carry information from the original message 20. The ACF 34 is related to controlling transmission of the slots through the network.

The TYPE field 36 within the slot can be used for the indication of a wide range of information. In one use it can allow for the separation of different communication classes within the one environment. As an example the messages with different field sizes and structures could be carried on the same switch. Also the TYPE field can be used to identify between different slot structures. Other segmentation schemes such as those described in the Background Art could be used in the same environment and distinguished by the TYPE field.

The TYPE field specifies how the SI field 38 and the information field of each slot should be interpreted depending on, among others, the following factors:

- (a) Whether connectionless or connection oriented communication is supported.
- (b) Whether short or long source and destination addresses are to be supported.
- (c) Whether extended hierarchical or non hierarchical source and destination addressing is supported.

09919765 073107

- (d) Whether the slot is carrying the first, last or continuation segment of a message.
- (e) Whether the message fits in a single segment.
- (f) Whether or not the source supports message transmission to multiple destinations concurrently.

The SI field 38 is a label which enables the logical association of all segments 40 belonging to the one message and hence enables them to be reassembled into the original message 20.

FIG. 2 diagrammatically illustrates a communications system comprising a source 42 which produces messages 20 of variable length for transmission on a network or fast packet switch 44 to a destination 46. The system includes a segmentation machine 48 coupled between the source 42 and the switch 44 and a reassembly machine 50 coupled between the switch 44 and the destination 46. The segmentation machine 48 converts messages 20 of indefinite length to slots 32 of fixed length for transmission on the switch 44. The reassembly machine 50 reassembles the slots 32 into the original message 20 for input to the destination 46. The segmentation and reassembly machines 48 and 50 would be located at respective nodes or access units coupled to the network.

FIG. 3 illustrates in more detail the use of the segmentation and reassembly machines 48 and 50 in a QPSX network of the type disclosed in WO 86/03639. The QPSX network comprises two unidirectional buses, bus A and bus B with data flowing in opposite directions, a central controller 2 and a number of distributed nodes or access units (AU's) 4 coupled between the buses A and B. Although each bus originates and terminates at the central controller 2 neither has a through connection, in normal circumstances. Each AU 4 has read taps 6 and 8 from the respective buses and lines 10 and 12 connected to unidirectional write couplers to the respective buses. The write transmit only in the direction of propagation of the respective buses. The read connections for each AU are attached to the bus ahead of the write connections and consequently the information read by each AU is unaffected by that written by it. In the illustrated arrangement, a source 42 is coupled to one of the access units 4 via the segmentation machine 48. The access unit transmits the message in fixed length slots on the network to the access unit 4 associated with the destination 46. Normally each access unit would have both segmentation and reassembly machines 48 and 50 to enable two way communications. The segmentation and reassembly machines 48 and 50 can be regarded as part of the interface IP 16.

FIG. 5 shows in more detail the connection of the segmentation and reassembly machines 48 and 50 to an access unit 4 of the type described in the aforementioned publication. FIG. 5 corresponds generally to FIG. 12 of that specification and hence need not be described in more detail here.

Each node in the network will have one or more unique SI's. Each SI can be used by the node for the transfer of a message. When the message transmission is complete the SI can be reused. Multiple SI's for a single node allow that node to transfer more than one message concurrently.

To describe the operation of the message transfer scheme, the segmentation of the message into slots is considered first and the action at the receiver is considered after that.

The train of slots 32 sent by the segmentation machine 48 is shown in FIG. 1. The first slot of a multiseg-

ment message will be identified as such by a COM (Beginning Of Message) code in the TYPE field 36. The SI field 36 is set by reference to the unique SI of the source node and the information field 40 contains the first segment of the message. Thus the DA field 22 of the message 20 is at the head of the information field 40. The following segments of the message until the last are each placed in the information fields of slots with the TYPE field 36 set to COM (Continuation Of Message) and the SI field 38 containing the source's code which is unique for this message. The last slot of a multisegment message has the TYPE field 36 coded as EOM (End Of Message), as shown.

For the transfer of a message 20 that only requires a single slot 32 the SSM (Single Slot Message) code is used in the TYPE field 36. The SI is not required in this case, however it is still used for consistency in operation.

An implementation of the segmentation machine 48 will now be described with reference to the state diagram which is shown in FIG. 6. In this diagram, the condition for a state transition is shown above the transition line and the action taken is below the line, in accordance with standard notation. This state machine will handle the receipt of a message from a single source 25 at a time. If simultaneous receipt of more than one message is required, multiple state machines would be required.

To specify the communications between adjacent parts of the system, there are three communications 30 primitives, as follows:

- (i) Request (REQ): This is a request to send a unit of data.
- (ii) Indication (IN): This is an indication that a unit of data has been received, and
- (iii) Confirmation (CONFIRM): This is a confirmation that a unit of data has been sent without error.

Between the machines 48 and 50 and the fast packet switch 44, (via the access units 4) communication primitives are prefixed by SAR\_ACC hence there are three 40 communication primitives as follows:

SAR\_ACC Request  
SAR\_ACC Indication  
SAR\_ACC Confirm.

Also for source or destination equipment 42 and 46 45 such as a computer attached to the segmentation and reassembly machines 48 and 50, communication primitives are prefixed by APP\_SAR hence there are three communication primitives as follows:

APP<sub>13</sub> SAR Request:  
APP\_SAR Indication  
APP\_SAR Confirmation

A segmentation machine operates on an unconfirmed APP SAR request, provided that sufficient resources are available. Tags for SAR\_ACC requests and Source 55 Identifiers (SIs) are common resources for all of these machines. Tags used in SAR\_ACC requests are unique over all such requests from any segmentation machines. A tag is allocated on an SAR\_ACC request and deallocated on an SAR\_ACC confirm. The tag for the 60 segments are local to the particular segment at the segmentation machine and is not transmitted on the network. Tags are reused when message transmissions have been completed. The tag is coded in TAG fields 52 and 54 in Request and Confirm primitives 56 and 58. 65 The TAG fields 52 and 54 in the Request and Confirm primitives 56 and 58 are used for communications between the segmentation machine 48 and access unit 4, as

09919725 "077101

see FIG. 7. The Request primitive 56 comprises the TAG field 52, T-SEG field 60 together with the slot 32. The fields 60 and 52 are control fields which are not transmitted beyond the access unit 4. The Confirm primitive 58 comprises a TAG field 54 and TR field 64, being control fields for communication between the access unit 4 and the segmentation machine 48. The codes in the TAG fields 52 and 54 must correspond for a particular segment 40. Therefore a particular message 20 will have a number of TAG codes corresponding to the various segments 32 required to be used for the transmission. Successful transmission or otherwise is indicated by the coding in the TR field 64.

Source Identifiers (SI) are selected at the source so that the SI will identify a unique message to any reassembly machine. The SI is allocated on an APP\_SAR request and is deallocated in an implementation dependent manner.

One reassembly machine 50 is needed at the destination 46 for each SI. A reassembly packet timer (not shown) is also associated with each reassembly machine, to prevent locking in a wait state for signals.

An SAR\_ACC confirm is routed to the segmentation machine which generated the SAR\_ACC request. An SAR ACC indication is routed to the reassembly machine associated with the SI. The segmentation function of the segmentation machine 50 is described in the following paragraphs.

#### 1.1 ABBREVIATIONS USED IN THE SEGMENTATION STATE DIAGRAM OF FIG. 6

35	CONFIRMSBACK (TAG)	CONFIRMSBACK is an array of flags indexed by tag number. A CONFIRMSBACK flag is set when the SARACC confirm has been received corresponding to the SARACC request with tag number TAG.
40	ACB:	All Confirms Back. Set if CONFIRMSBACK(1) is set for all segments.
	CHECKSUM:	Flag for checking sum of bits successfully transmitted.
45	LAST SEGMENT SENT:	Flag indicating whether all segments for this message has been sent.
	LAST SEGMENT SENT:	Flag indicating that not all segments for this message has been sent.
50	SI:	The source identifier from the segment in an SARACC indication.
	TR:	The TR field 64 from the SARACC confirm indicates SUCCESS or FAILURE in the transmission of a segment.
55	TSEG:	The TSEG field 60 in the SAR ACC request indicates whether the segment is the first segment of the message (BOM), a continuing segment (COM), or the final segment (EOM).

The states of the state machine diagram are as follows:

#### 1.2 STATE S0:IDLE

State S0 is the initial state for all segmentation state machines. In this state, no APP\_SAR request is outstanding from the source 42.





T\_SEG: The T\_SEG field in the SAR\_ACC Indication indicates whether the slot 32 is the first segment of the message (BOM), a continuing segment (COM) or the final segment (EOM).

## 2.2 STATE R0: IDLE

R(00a) Idle—An SAR\_ACC indication presents a segment with SSM set. A single segment message is assembled. The checksum is computed and if no error is indicated, an APP\_SAR Indication presents the complete message to the destination 46.

R(01) Receiving—An SAR—ACC Indication presents a segment with BOM set and sufficient SAR resources are available to allow the reassembly machine to leave the idle state. Buffer allocated of size in octets equal to LENGTH indication. Resources are allocated and the data from the segment is buffered.

In FIG. 8, Action 1 beneath the transition line (11) includes the steps of storing the received segment 32 in sequence and computing the CHECKSUM over the message 20. Action 2 in transition lines R0(00a) and R1(10b) includes of reassembling the message 20, computing the CHECKSUM over the message 20. If CHECKSUM IN GOOD then APP\_SAR\_IND, otherwise discard the message.

In this state, the reassembly state machine is reassembling a message. The machine will remain in this state 45 until the complete message has been assembled or until the Reassembly Packet Timer PRT has expired or there is a problem with available resources. The RPT ensures that the reassembly machine is not effectively put out of service waiting for a lost EOM.

R(10b) Idle—When the complete message has been assembled (EOM received) the reassembly machine computes the CHECKSUM. If the CHECKSUM indicates no error, an APP\_SAR indication presents the received message to the destination 46. If the CHECKSUM indicates an error, the reassembly machine returns to the Idle state and releases resources.

The reassemble state machine diagram shown in FIG. 8 could be implemented in logic or by appropriate software.

FIG. 9 is another example of a state machine diagram for a reassembly machine which is particularly designed for use with the distributed queue system of the OPSX

network disclosed in the aforementioned international publication.

This receive machine has two states: IDLE and WAIT. In the IDLE state the machine is not currently receiving any message. Thus in this state the machine will check for slots 32 with TYPE field 36 equal to BOM or SSM. In the case a BOM code is received the machine will check the DA in the information field. If the message is addressed to the station the machine enters the second state. In the case that SSM is detected in the IDLE state and the DA field matches, the length and information fields are copied and the AU 4 will indicate the higher layers that a message is received. The receive machine remains in the IDLE state after copying the slot. The higher layers refer to the higher layers in the Open System Interconnections (OSI) reference model.

The WAIT state is used to receive the slots 32 following the first of a multisegment message. In this state the slots with TYPE field 36 equal to COM and the SI equal to that copied from the first slot of the message will be received. The information fields of these following slots are concatenated to form the complete message. New messages addressed to the given station in this state are ignored by the state machine. Further receive machines are required if such messages are to be received. When the last slot of the message is detected by the receive state machine, the machine will copy the information field of the slot, indicating to the higher layers that a message is received and return to the IDLE state. This completes the receipt of the message. To guard against the loss of the EOM slot, which would cause the receive state machine to be locked in the WAIT state, the RPT timer is used. This timer is started after each BOM or COM slot is received. If the timer expires before the next slot is received, a failure in the transfer is assumed. The machine will then clear all copied slots and return to the IDLE state.

#### CONSTRAINED DESTINATION RESOURCES 40

In all practical implementations of destination facilities, i.e. resources such as buffers and processing capacity will be limited and loss of slots can occur unless a positive control mechanism is introduced to control communication between source and destination.

The method enables a source seeking use of the destination resources for reassembly of the original message to be temporarily held up until the required resources are available. The method implements a variety of access disciplines. Among others, the method supports first come first served queueing for the destination resources described briefly below.

Control of access to destination resources is via a "ticket" handout scheme in which the "ticket number" establishes the position of source request in a distributed queue. The first encapsulated segment of a message sent by a source is considered by the destination as a request for resources. A "go ahead" message is returned to the source if the required resources are available, otherwise a "ticket" which uniquely identifies the sources position in a distributed queue of requests is returned. When resources become available the destination broadcasts the "ticket number" of the next source to be serviced. Under the assumption that resources are usually available, the delay incurred in waiting for the "go-ahead" is avoided by the source continuing to send segments of a message while it awaits the reply from the destination. In the case of a negative reply, that is a "ticket" is re-

099165 073703

—

3

10

30

55

65

FIGS. 12a and b are flow charts which illustrate the logical steps involved in the reassemble machine 50, illustrated in FIG. 11.